

Computer Model Calibration with Large Nonstationary Spatial Outputs

KAI-LAN CHANG AND SERGE GUILLAS*

University College London
ucakkac@ucl.ac.uk

Abstract

The Bayesian computer model calibration method has proven to be effective in a wide range of applications. In this framework, input parameters are tuned by comparing model outputs to observations. However, this methodology becomes computationally expensive for large spatial model outputs. To overcome this challenge, we employ a truncated basis representations of the model outputs. We then aim to match the model outputs coefficients with the coefficients from observations in the basis representations; we also optimize the truncation level. In a second step, we enhance the calibration with the addition of the INLA-SPDE technique. We incorporate the nonstationary behavior and the derivative information of the spatial field into the calibration by inserting two INLA-SPDE parameters into the calibration. Several synthetic examples and a climate model illustration highlight the benefits of our approach for model outputs distributed over the plane or the sphere.

Keywords: dimension reduction; SPDE; Matérn fields; uncertainty quantification

1 Introduction

Complex computer models are widely used in various fields of science and technology to mimic complex physical systems. Computer model calibration involves comparing the simulations of a complex computer model with the physical observations of the process being simulated. Increasingly, computer model outputs are in the form of spatial fields, particularly in environmental sciences. This poses a particular challenge to the calibration method. In this paper we develop our Bayesian calibration technique based on the framework from Kennedy and O’Hagan (2001): we approximate the expensive computer model by a Gaussian process (GP). This formulation has proven to be effective in a wide range of applications. However, the GP calibration is computationally expensive for large model input and output spaces (typically cubic in the number of data points used to fit the GP). Therefore several attempts to tackle this issue have been made by using truncated basis representations of model outputs in order to reduce dimension (Bayarri et al., 2007; Chang et al., 2014; Higdon et al., 2008), by focusing on the plausibility of model outputs v. the observations in order to constrain the input spaces (McNeall et al., 2013; Williamson et al., 2013), or by using a separable covariance function over space and tuning parameters to build a theoretical emulator for multivariate outputs (Bhat et al., 2010; Rougier, 2008). We provide here a solution that makes use of an adequate representation of the spatial outputs using Gaussian fields.

Gaussian fields (GF) play an important role in spatial statistics and form a major area of interest within the field of Bayesian hierarchical spatial models. The traditional approach is to specify a GF through its covariance function. Another approach is to use the class of Gaussian Markov random fields (GMRF), which are discretely indexed GFs. The Markov property yields a sparse precision matrix, so that efficient numerical algorithms can be employed. Lindgren

*Kai-Lan Chang was supported by the Taiwanese government sponsorship for Ph.D. overseas study. S. Guillas was partially supported by a Leverhulme Trust research fellowship on “stratospheric ozone and climate change” (RF/9/RFG/2010/0427). We also thank Hanli Liu (NCAR) for technical support in running the chemistry-climate model WACCM.

et al. (2011) show that the GMRF representation can be constructed explicitly by using a certain form of stochastic partial differential equation (SPDE) which has a GF with Matérn covariance as its solution when driven by Gaussian white noise. The representation employs piecewise linear basis functions, and Gaussian weights with Markov dependences determined by the finite element method over a triangulation of the domain. This technique can deal with large spatial data sets and naturally account for nonstationarity. Our paper combines both the calibration formulation and the SPDE defined scale and precision parameterization to deal with large scale spatial outputs, and still provides a compromise with computational feasibility in order to employ a fully Bayesian approach.

To describe the general framework, let $\eta(\mathbf{s}_i, \boldsymbol{\theta}_j), i = 1, \dots, n; j = 1, \dots, r$ be the r -runs model output measured at n locations. Here we refer to $m = n \times r$ as the total number of outputs in the simulations. We choose a design made of combinations of input values, and we impose distributional prior assumptions on the inputs $\boldsymbol{\theta}$. The aim of calibration is to compare outputs to observations and estimate the best input setting $\boldsymbol{\theta}^*$. Our practical interest in this paper is to tune, and quantify uncertainty in, climate experiments. The Whole Atmosphere Community Climate Model (WACCM) is a general circulation model of the middle and upper atmosphere. WACCM is an extension of the National Center for Atmospheric Research Community Earth System Model (CESM). Many parameterizations of physical processes have to be set to run WACCM, resulting in potential concerns about error growth (Liu et al., 2009). Specifically, the gravity wave (GW) parametrization in climate models aims to reduce zonal mean wind biases. Small modification of parameterized GWs can have large impacts by improving the propagation pathways of the Rossby waves (Alexander and Sato, 2015). GWs also play a dominant role in driving the quasi-biennial oscillation (QBO) in the Tropical stratosphere. For this study we explore the zonal wind outputs over the sphere in four GW inputs with $r = 100$ runs in order to calibrate the GW parameters. Each spatial output was computed by the WACCM simulator over a grid of $n = 96 \times 144 = 13824$ cells. The resulting number of computer runs r , combined with the output dimension n may become too large to fit GPs to the computer model, and the discrepancy between model and observations, and thus may prevent the fully Bayesian calibration to be performed.

In Section 2.1 we present our approach in detail. We employ a truncated basis representation, such as B-splines decomposition, Fourier or spherical harmonics transform, to capture the output features spatially. We then explore how parameters in an SPDE model can explicitly quantify the nonstationarity of the spatial field (Bolin and Lindgren, 2011; Zammit-Mangion et al., 2015). Hence in Section 2.2 we extend our approach by including the basis representation of the scale and precision parameters in an SPDE model into our calibration framework. In Section 3, we apply these techniques to different synthetic examples and our real climate experiment, with the aim of evaluating performance of our calibration method under different sources of uncertainty. Finally in Section 4, we discuss potential improvements to our approach.

2 Methods

2.1 Basis representation for the model output

In this section, we derive a basis representation for surface model outputs and observations using real valued basis functions, such as B-splines, spherical harmonics, or some other bases, to represent these surfaces, and construct a methodology for the calibration employing the coefficients in these representations. Let $\boldsymbol{\theta}$ be the calibration parameters. The output $\eta(\cdot)$ is computed at inputs $(\mathbf{s}, \boldsymbol{\theta})$ in an m -point experimental design, where $m = n \times r$ means r computer runs measured at n locations. The output $\eta(\mathbf{s}, \boldsymbol{\theta})$ is an approximation of the reality $y^R(\mathbf{s})$. The discrepancy between the simulator and the reality at the spatial locations is denoted $\delta(\mathbf{s})$. The observations $y^F(\mathbf{s})$ of reality are collected at a number of locations \mathbf{s} in an

n -point spatial design (here a simple grid), and are subject to an observation error $\epsilon(\mathbf{s})$. The measurement locations for observations and outputs can be different, since the methodology accommodates such variation. According to Kennedy and O'Hagan (2001),

$$y^F(\mathbf{s}) = y^R(\mathbf{s}) + \epsilon(\mathbf{s}) = \eta(\mathbf{s}, \boldsymbol{\theta}^*) + \delta(\mathbf{s}) + \epsilon(\mathbf{s}).$$

We use a set of basis functions $\{\psi_z\}$, where z is an integer that represents the index number within the ordered basis, to decompose each run of model output. Precisely, for the N_η -th level of expansion and for each run j :

$$\eta(\mathbf{s}, \boldsymbol{\theta}_j) = \sum_{z=1}^{N_\eta} c_z^M(\boldsymbol{\theta}_j) \psi_z(\mathbf{s}) + \varepsilon_{N_\eta}, \quad j = 1, \dots, r,$$

where ε_{N_η} is an error term associated with the expansion order N_η . The coefficients $\{c_z^M\}$ represent the surface frequency or curve features at different resolutions. This representation obviously separates the spatial effect from the variations triggered by the inputs $\boldsymbol{\theta}$. The variations in $\boldsymbol{\theta}$ propagate only to the expanded coefficients. Similarly, the observations can be written as

$$y^F(\mathbf{s}) = \sum_{z=1}^{N_y} c_z^F \psi_z(\mathbf{s}).$$

By applying a basis representation, the physical space of both model outputs and observations are transformed into a functional space. Since the aim is to calibrate the spatial outputs, we also assume that the reality $y^R(\mathbf{s})$, the discrepancy function $\delta(\mathbf{s})$, and the measurement errors $\epsilon(\mathbf{s})$, can be represented by similar basis representations:

$$y^R(\mathbf{s}) = \sum_{z=1}^{N_y} c_z^R \psi_z(\mathbf{s}), \quad \delta(\mathbf{s}) = \sum_{z=1}^{N_y} c_z^\delta \psi_z(\mathbf{s}), \quad \epsilon(\mathbf{s}) = \sum_{z=1}^{N_y} c_z^\epsilon \psi_z(\mathbf{s}).$$

Since the computer model is implemented through a complex mathematical code, the spatial outputs from the computer simulation should be relatively smoother than the observations. Therefore we assume a larger number of basis functions (N_y) in the observations (as well as discrepancy and error functions) than for model outputs (N_η , $N_\eta \leq N_y$). In the formulation of the calibration algorithm, all coefficients $\{c_z^M | N_\eta < z \leq N_y\}$ are set to be 0 in order to obtain a sparse correlation matrix and maintain the computational feasibility. To simplify the notation however, we use the same number of basis N_y to decompose y^F and η , but different levels of expansion would work as well as long as these are larger than N_η . Then matching the coefficients yields

$$c_z^F = c_z^R + c_z^\epsilon = c_z^M(\boldsymbol{\theta}^*) + c_z^\delta + c_z^\epsilon, \quad z = 1, \dots, N_y.$$

Hence, only the relatively smooth variations of the computer model match the variations in observations. A careful choice of N_η and N_y allows an optimal comparison. The parameters of the measurement errors, c_z^ϵ , are assumed normally distributed with mean 0 and variance σ_z^2 , and independent across z . The coefficients $\{c_z^M\}$ can be assumed to have mean 0 with marginal variance close to 1 without loss of generality, if the responses are standardized. The

GP assumption is imposed on each coefficient $c_z^M(\boldsymbol{\theta})$, $z = 1, \dots, N_y$, of mean 0 and with a covariance function

$$\text{Cov}(c_z^M(\boldsymbol{\theta}), c_{z'}^M(\boldsymbol{\theta}')) = \frac{1}{\lambda_\eta} I_{zz'} \prod_{k=1}^q \rho_{\eta k}^{4(\theta_k - \theta'_k)^2}, \quad (1)$$

where $I_{zz'}$ is the Kronecker's delta ($I_{zz'} = 1$ if $z = z'$ and 0 otherwise), q is the dimension of $\boldsymbol{\theta}$, λ_η controls the marginal precision of $\eta(\cdot, \cdot)$ and ρ_η controls the strength of the dependence in each of the pairs of $\boldsymbol{\theta}$. This formulation leads to a smooth and infinitely differentiable representation for the model output. In addition, a set of coefficients decomposed by the same basis ψ_z forms a block in the covariance structure: we assume that the correlation between different indices z of the coefficients is 0. Hence the rN_y -vector \mathbf{c}^M has a multivariate normal prior with mean 0 and a covariance matrix with $r \times r$'s N_y blocks in the diagonal and the off-diagonal blocks are zero matrices. The specifications for λ_η and ρ_η are gamma priors and independent beta priors: $\pi(\lambda_\eta) \sim \lambda_\eta^{a_\eta - 1} e^{-b_\eta \lambda_\eta}$ and $\pi(\rho_{\eta k}) \sim \rho_{\eta k}^{a_{\rho_\eta} - 1} (1 - \rho_{\eta k})^{b_{\rho_\eta} - 1}$, $k = 1, \dots, q$.

We assume that the discrepancy (also called bias) $\delta(s)$ in the physical process displays a spatial pattern that can be represented and reconstructed by a set of basis functions. The decomposed discrepancy term c_z^δ quantifies the inadequacy between the simulator and reality in the functional domain, which is independent of the c_z^M . We assume that c_z^δ follows a GP model:

$$c_z^\delta \sim N\left(\mathbf{0}, \frac{1}{\lambda_\delta} I_{zz'}\right). \quad (2)$$

The level of this expansion must be chosen carefully as the truncated basis may not have enough degrees of freedom to capture all of the variability in the observations (typically much more variability than in the smooth model outputs). The prior specification for the parameter λ_δ is a gamma prior $\pi(\lambda_\delta) \sim \lambda_\delta^{a_\delta - 1} e^{-b_\delta \lambda_\delta}$.

Denote the joint $(r + 1)N_y$ data vector $\mathbf{D} = (c^F, c^M)$. The sampling likelihood for the full data is then

$$L(\mathbf{D} | \boldsymbol{\theta}, \lambda_\eta, \rho_\eta, \lambda_\delta, \Sigma_y) \propto |\Sigma_{\mathbf{D}}|^{-1/2} \exp\left\{-\frac{1}{2}(\mathbf{D}^T \Sigma_{\mathbf{D}}^{-1} \mathbf{D})\right\}, \quad (3)$$

where

$$\Sigma_{\mathbf{D}} = \Sigma_\eta + \begin{pmatrix} \Sigma_\epsilon + \Sigma_\delta & 0 \\ 0 & 0 \end{pmatrix},$$

in which Σ_ϵ is the $N_y \times N_y$ observation covariance matrix, Σ_η is obtained for each pair of $(r + 1)N_y$ simulation inputs through (1) corresponding to \mathbf{D} , and Σ_δ is an $N_y \times N_y$ matrix obtained for each pair of N_y input through the instances of (2) that correspond to the coefficients c^F .

Let $\pi(\boldsymbol{\theta})$ be the joint prior distribution for the (unknown) calibration vector $\boldsymbol{\theta}$. The resulting posterior density has the form

$$\pi(\boldsymbol{\theta}, \lambda_\eta, \rho_\eta, \lambda_\delta | \mathbf{D}) \propto L(\mathbf{D} | \boldsymbol{\theta}, \lambda_\eta, \rho_\eta, \lambda_\delta) \times \pi(\boldsymbol{\theta}) \times \pi(\lambda_\eta) \times \pi(\rho_\eta) \times \pi(\lambda_\delta), \quad (4)$$

which can be explored via a Markov chain Monte Carlo (MCMC) technique, for which we employ a Metropolis–Hastings algorithm. The calibrated vector is then denoted by $\boldsymbol{\theta}^* = \text{argmax}_{\boldsymbol{\theta}} \pi(\boldsymbol{\theta}, \lambda_\eta, \rho_\eta, \lambda_\delta | \mathbf{D})$.

The strong assumption of independence of the coefficients, through different blocks in the covariance, may not be fully justifiable in real applications. However, this assumption leads to a great computational advantage in terms of forming a block diagonal covariance model in the GP model. Traditionally a GP fitting involves a computational complexity of $O(m^3) = O(n^3 r^3)$. In the basis representation approach the complexity of our model is $O(N_y^3 r^3)$, where $N_y \ll n$. The block diagonal assumption further reduces the cost to $O(N_y r^3)$. In section 3 we discuss how this assumption is a compromise between fidelity and complexity. Furthermore, it is

not common knowledge to know how to specify the covariance structures for the truncated basis representation due to the mathematical challenge of finding explicit expressions for the covariance (Jun and Stein, 2008). Nevertheless, there is an alternative way to efficiently model complex spatial covariance structures with the added bonus of a suitable depiction of the nonstationarity structure into our calibration algorithm. We introduce it in the next section.

2.2 Calibration with SPDE modeling

Models in spatial statistics that enable to build an approximation of the entire underlying random field are usually specified through the covariance function of the latent field. In order to assess uncertainties in the spatial interpolation over the whole spatial domain, we cannot build models only for the discretely located observations or model outputs, we need to build an approximation of the entire underlying stochastic process defined on the spatial field. We consider statistical models for which the unknown functions are assumed to be realizations of a Gaussian random spatial process. The traditional fitting approach spatially interpolates values as linear combinations of the original observed locations, and this constitutes the spatial kriging predictor. Not only large data sets can be computationally demanding for a covariance-based approach, but this approach struggles to take into consideration nonstationarity, due to the fixed underlying covariance structure. A different computational approach was introduced by Lindgren et al. (2011), in which random fields are expressed as a weak solution to an SPDE, with explicit links between the parameters of the SPDE model and the Matérn covariance function. In this section we review some of the main concepts in spatial modeling through the SPDE approach. We describe the details of our extension by including the SPDE defined scale and precision parameters into the Bayesian calibration.

2.2.1 Spatial modeling through the SPDE approach

The Matérn covariance function is an advanced covariance structure used to model dependence of spatial data on the plane. The shape parameter ν , the scale parameter κ , and the marginal precision τ^2 , parameterize it:

$$\text{Cov}(\mathbf{h}) = \frac{2^{1-\nu}}{(4\pi)^{d/2} \Gamma(\nu + d/2) \kappa^{2\nu} \tau^2} (\kappa \|\mathbf{h}\|)^\nu K_\nu(\kappa \|\mathbf{h}\|), \mathbf{h} \in \mathbb{R}^d,$$

where \mathbf{h} denotes the difference between any two locations s and s' : $\mathbf{h} = s - s'$, and K_ν is the modified Bessel function of the second kind of order $\nu > 0$.

We denote by $Y(\mathbf{s})$ the observations (or the spatially distributed model outputs) for a latent spatial field $X(\mathbf{s})$, with a Matérn covariance structure. We assume a zero mean Gaussian noise, $\mathcal{W}(\mathbf{s})$, with a constant variance σ_s^2 : $Y(\mathbf{s}) = X(\mathbf{s}) + \mathcal{W}(\mathbf{s})$. Thus, according to Whittle (1963), the latent field $X(\mathbf{s})$ is the solution of a stationary SPDE:

$$(\kappa^2 - \Delta)^{\alpha/2} \tau X(\mathbf{s}) = \mathcal{W}(\mathbf{s}). \quad (5)$$

where Δ is the Laplace operator. We explain in the next paragraph how the analysis of this SPDE can be carried out by the finite element method. The regularity (or smoothness) parameter ν essentially determines the order of differentiability of the fields. The link between the Matérn field and the SPDE is given by $\alpha = \nu + d/2$, which makes explicit the relationship between dimension and regularity for fixed α . On more general manifolds than \mathbb{R}^d , such as the sphere, the direct Matérn representation is not easy to implement, but the SPDE formulation provides a natural generalization, and the ν -parameter will keep its meaning as the quantitative measure of regularity. Instead of defining Matérn fields by the covariance function, Lindgren et al. (2011) used the solution of the SPDE as a definition, and it is much easier and flexible to do so. This definition also facilitates nonstationary extensions by allowing the SPDE parameters κ and τ in Eq. (5) to vary with space, hence denoted $\kappa(\cdot)$ and $\tau(\cdot)$ respectively.

We estimate the SPDE parameters and supply uncertainty information about the spatial fields by using the *integrated nested Laplacian approximations* (INLA) framework, available as an R package (<http://www.r-inla.org/>) (Lindgren and Rue, 2015). The models implemented in the INLA–SPDE framework are built on a basis representation (triangulation over the spatial domain): $X(\mathbf{s}) = \sum_{i=1}^M \varphi_i(\mathbf{s})w_i$, where $\{w_i\}$ are the stochastic weights chosen so that the distribution of the functions $X(\mathbf{s})$ approximates the distribution of solutions to the SPDE on the space, and $\varphi_i(\mathbf{s})$ are piecewise linear basis with compact support (i.e. finite elements) in order to obtain a Markov structure, and to preserve it when conditioning on local observed locations. The Markov property yields a sparse precision matrix, so that efficient numerical algorithms can be employed for large spatial data. The projection of the SPDE onto the basis representation is chosen by a finite element method. The finite element method represents a general class of techniques for the approximate solution to boundary value problems for partial differential equations. The piecewise linear basis functions defined by a triangulation of the spatial domain allow us to explicitly evaluate the precision matrix of the latent field. As a result, $X(\mathbf{s})$ follows a normal distribution with mean 0, and the precision matrix can be explicitly expressed as a combination of the piecewise linear basis functions weighted by κ and τ (which means κ and τ have a joint influence on the marginal variances of the latent field). Then $X(\mathbf{s})$ can be generated continuously as approximative solutions to the SPDE.

Note that the INLA–SPDE model is constructed through a triangulation. The triangulation can be made adaptive to the irregularly distributed spatial data. Here our application uses regularly gridded spatial outputs and observations. It is beyond the scope of this paper to discuss flexible and irregular triangulation (though example 1 shows an illustration). The default value in INLA is $\alpha = 2$, but $0 \leq \alpha < 2$ are also available, though yet to be completely tested (Lindgren and Rue, 2015). So with a 2-manifold (both in \mathbb{R}^2 and \mathbb{S}^2), the smoothness parameter ν must be fixed at 1 due to the relationship $\alpha = \nu + d/2$. For the prior specifications of the INLA–SPDE model, κ and τ follow log normal priors by default. We do not change the default prior settings throughout the analysis.

2.2.2 Combining SPDE modeling and calibration

To allow a nonstationary extension of the INLA–SPDE model (with a focus on the calibration), let $\kappa^M(\mathbf{s}, \boldsymbol{\theta})$ and $\tau^M(\mathbf{s}, \boldsymbol{\theta})$ be the scale and precision parameters in an SPDE model used to approximate computer model outputs. To obtain basic identifiability, $\kappa^M(\mathbf{s}, \boldsymbol{\theta})$ and $\tau^M(\mathbf{s}, \boldsymbol{\theta})$ are taken to be positive, and their logarithm can be decomposed as:

$$\log \kappa^M(\mathbf{s}, \boldsymbol{\theta}_j) = \sum_{z=1}^{N_\kappa} \kappa_z^M(\boldsymbol{\theta}_j) \psi_z(\mathbf{s}), \quad \text{and} \quad \log \tau^M(\mathbf{s}, \boldsymbol{\theta}_j) = \sum_{z=1}^{N_\tau} \tau_z^M(\boldsymbol{\theta}_j) \psi_z(\mathbf{s}), \quad j = 1, \dots, r.$$

We use a basis representation for both the scale and precision parameters in a nonstationary SPDE model. As $\kappa_z^M(\boldsymbol{\theta})$ and $\tau_z^M(\boldsymbol{\theta})$ can quantify the nonstationarity and derivative information in the spatial process, we now include these two coefficients into our technique (combined with $c_z^M(\boldsymbol{\theta})$ in the previous section). Then the model outputs consist of 3 type of coefficients $\{c_{z,j}^M, \kappa_{z,j}^M, \tau_{z,j}^M\}, j = 1, \dots, r$ in associated with the model discrepancy previously presented. Indeed, we represent the spatial nonstationarity in the observations with coefficients $\{\kappa_z^F, \tau_z^F\}$. Then the observations consist of the coefficients $\{c_z^F, \kappa_z^F, \tau_z^F\}$.

The aim is to combine the scale and precision parameters as nonstationary covariance information for the implementation of the calibration algorithm, and to model all coefficients jointly with the GP assumption. The expansion order of the basis functions for $\{c_{z,j}^M, \kappa_{z,j}^M, \tau_{z,j}^M\}$, do not need to be the same, and coefficients with different indices z are treated as independent. We also assume that the three types of coefficients are independent. To describe the formulation of the design matrix, let $\{z_1, z_2, z_3 | z_1 = 1, \dots, N_y; z_2 = 1, \dots, N_\kappa; z_3 = 1, \dots, N_\tau\}$ be the number

of coefficients used to represent each triplet of coefficients, respectively. Thus there are $(N_y + N_\kappa + N_\tau)$ -blocks of coefficients corresponding to each combination of $\theta_j, j = 1 \dots, r$ in the covariance matrix. The GP assumption is imposed on each coefficient $(c_{z_1,j}^M, \kappa_{z_2,j}^M, \tau_{z_3,j}^M)^T$ with mean 0 and covariance function

$$\text{Cov}((c_{z_1}^M(\theta), \kappa_{z_2}^M(\theta), \tau_{z_3}^M(\theta)^T), (c_{z_1'}^M(\theta'), \kappa_{z_2'}^M(\theta'), \tau_{z_3'}^M(\theta')^T) = \frac{1}{\lambda_\eta} \prod_{i=1}^3 I_{z_i z_i'} \prod_{k=1}^q \rho_{\eta k}^{4(\theta_k - \theta_k')^2}.$$

In other words, these 3 types of coefficients $\{c_{z_1}^M, \kappa_{z_2}^M, \tau_{z_3}^M\}$ have a joint multivariate normal prior distribution with mean 0, with covariance structure forms a block diagonal matrix:

$$\begin{pmatrix} c_{z_1}^M \\ \kappa_{z_2}^M \\ \tau_{z_3}^M \end{pmatrix} \sim N \left(\mathbf{0}, \begin{pmatrix} \text{Cov}(c_{z_1}^M(\theta), c_{z_1'}^M(\theta')) & 0 & 0 \\ 0 & \text{Cov}(\kappa_{z_2}^M(\theta), \kappa_{z_2'}^M(\theta')) & 0 \\ 0 & 0 & \text{Cov}(\tau_{z_3}^M(\theta), \tau_{z_3'}^M(\theta')) \end{pmatrix} \right).$$

The elements in each block are also block diagonal matrices. As we assumed independence between different blocks of coefficients, the model discrepancy term in the functional space follows a GP assumption defined in Equation (2). All the prior assumptions discussed in the previous section remain unchanged. Thus the sampling likelihood in (3) and the posterior distribution in (4) still hold in this case. Overall, we decompose the model outputs into a basis via the coefficients c^M , and estimate the SPDE parameters $\{\kappa^M, \tau^M\}$ in the latent field through a regression onto these basis functions. We are essentially fitting a GP model with c^M for the regression mean structure and $\{\kappa^M, \tau^M\}$ for the parameters of Matérn covariance function (with ν fixed in connection with SPDE model).

3 Simulation Study and Application

In order to illustrate the methodology, we present four synthetic examples and one real computer experiment. In example 1, we start with a simulation over irregularly spaced outputs over the plane as it is less complex than over the sphere. Then we extend our technique step-by-step to a spherical output. Example 2 shows how the dimension reduction works from a spherical output to a spherical representation. We investigate the connection between the calibration accuracy and the number of computer runs r , and between the calibration accuracy and the orders and modes of spherical harmonics (SHs). Example 3 is designed to compare different basis representations for the model output, with the aim of evaluating the respective performances of the SHs approach and the principal components (PCs) approach (Higdon et al., 2008). Example 4 simulates a nonstationary field on the sphere, with an anisotropic property (the spatial correlation depends on latitude), to demonstrate how including the parameters in the SPDE can further benefit the GP calibration. Finally, in our application to real data, we perform the calibration of gravity wave (GW) parameters for the WACCM simulator, with both known synthetic (but with added nonstationary observation errors) and real observations in order to fully validate our approach.

3.1 Example 1: Irregular grids

To illustrate our methodology of dimension reduction for the spatial output, this example is devoted to a spatial calibration using irregularly spaced locations in \mathbb{R}^2 . The true target function (i.e. computer model) is inspired by a bivariate normal function, defined as

$$f(\mathbf{x}, \theta) = \frac{1}{2\pi\theta_1\theta_2} \exp \left(-\frac{1}{2\pi} \left(\frac{(x_1 - 0.5)^2}{\theta_1^2} - \frac{(x_1 - 0.5)(x_2 - 0.5)}{\theta_1\theta_2} + \frac{(x_2 - 0.5)^2}{\theta_2^2} \right) \right), \quad (6)$$

with the true values for (θ_1, θ_2) being set to $(0.3, 0.7)$, and $\mathbf{x} = (x_1, x_2) \in [0, 1]^2$. Both calibration and location parameters are generated by a maximin Latin hypercube design (LHD). We generated noisy observations $y^F(\mathbf{x}, \boldsymbol{\theta}^*) = f(\mathbf{x}, \boldsymbol{\theta}^*) + \epsilon(\mathbf{x})$ at locations $(x_{1i}, x_{2i}), i = 1, \dots, 100$, on an irregularly spaced grid, where $\epsilon(\mathbf{x}) \sim N(0, 0.01)$. The model output $\eta(\mathbf{x}, \boldsymbol{\theta}) = f(\mathbf{x}, \boldsymbol{\theta}) + \delta(\mathbf{x})$ was generated employing the same 100 spatial locations as observations and 50 runs of a maximin LHD $\boldsymbol{\theta}_j = (\theta_{1j}, \theta_{2j}), j = 1, \dots, 50$, where $\delta(\mathbf{x}) \sim N(\frac{1}{2}\sigma_y^2, 0.01)$, with σ_y^2 the variance of the observational data y^F . Fig. 1(a) shows the distribution of the irregularly spaced grid (circles) in our design.

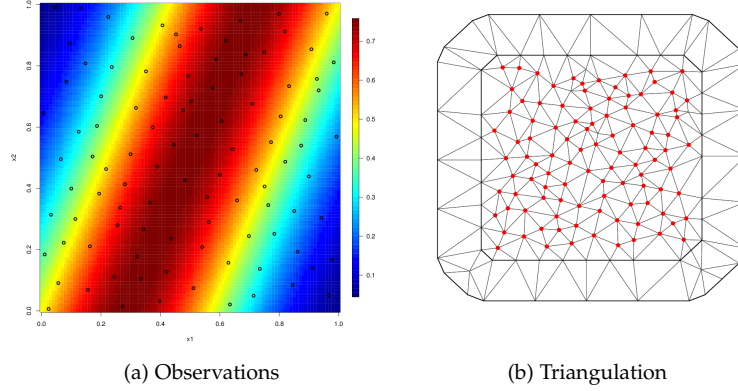


Figure 1: (a) Observations at 100 spatial locations (circles); (b) The spatial triangulation.

In order to investigate the benefit of incorporation of SPDE parameters into the calibration algorithm, the following models are compared: (1) Model A: 10 B-spline basis decomposition for η and y^F ; (2) Model B: 15 B-spline basis decomposition for η and y^F ; (3) Model C: 10 B-spline basis decomposition, including 2nd order B-spline for SPDE defined precision and scale parameters, for η and y^F . Fig. 1(b) shows the triangulation covering the spatial domain.

The specifications for the rest of hyperpriors are $\lambda_\epsilon \sim \Gamma(1, 0.003)$, $\lambda_\delta \sim \Gamma(1, 0.01)$, $\lambda_\eta \sim \Gamma(5, 5)$, and $\rho_{\eta k} \sim \text{beta}(1, 0.1), k = 1, \dots, 3$. Besides, to represent vague prior information about the true values of the calibration parameters, we specify a uniform prior for each component of $\boldsymbol{\theta}$ over $[0, 1]$ (all the prior specifications are the same throughout the paper).




Figure 2: The posterior densities of θ_1 (a) and θ_2 (b) for Model A (dotted line), B (dashed line) and C (solid line). Vertical lines indicate the true values for θ_1 and θ_2 .

The posterior distributions of θ_1 and θ_2 are shown in Fig. 2(a) and (b). The effect of including

SPDE parameters is clearly reflected in a decrease in the bias and the uncertainty in the posterior densities. For Model B, the posterior accuracy of θ_2 is better than Model A, as expected, when the size of the bases of functions is increasing. However, the precision in posterior density is only slightly improved. The posterior density of θ_1 shows a similar patterns between Model A and B. Model C uses fewer coefficients than Model B, while Model C achieves better results in both posterior accuracy and precision. This finding suggests that the calibration can be improved, along with parsimony, by incorporating SPDE parameters in the analysis.

3.2 Example 2: Number of experiments and size of expansion

Covariance models for data distributed on the sphere introduce additional constraints in contrast to covariance models for data over an Euclidean space. For example, the issues of choice of distance (great circle or chordal), or the property of positive definiteness of covariance functions on the sphere (Chang et al., 2015a; Gneiting, 2013). Note that Kennedy and O'Hagan (2001) showed that the GP calibration is relatively robust to different covariance structures, so the impact may be modest in some cases.

SHs constitute a frequency-space basis of functions defined on the sphere, which are equivalent to the Fourier series on the spherical domain. Denoting a location over an unit sphere $s = (L, l) \in S^2$ using latitude $L \in [0, \pi]$ and longitude $l \in [0, 2\pi)$, the real SHs function on an unit sphere $\psi_{k,h} : S^2 \rightarrow \mathbb{R}$ is defined as

$$\psi_{k,h}(s) = \sqrt{\frac{2k+1}{4\pi} \frac{(k-|h|)!}{(k+|h|)!}} \begin{cases} \sqrt{2} \sin(hl) P_{k,|h|}(\cos L) & \text{if } -k \leq h < 0, \\ P_{k,0}(\cos L) & \text{if } h = 0, \\ \sqrt{2} \cos(hl) P_{k,h}(\cos L) & \text{if } 0 < h \leq k, \end{cases} \quad (7)$$

where $P_{k,h}$ is the associated Legendre polynomials, the order k determines the frequency of the basis functions over the sphere, and the mode h gives the longitudinal frequency and is generally called the wavenumber in the geophysical literature. Furthermore, the Fourier transform of a GP is still Gaussian (Kirby, 2001). The SHs form a complete set of orthonormal functions in S^2 . On the sphere, any square integrable function can thus be expanded as a linear combination of SHs, so SHs are a natural way to represent a function on the sphere.

The second example illustrates the efficiency of the SHs transform (SHT) for a large spherical output. Given an element $\mathbf{s} = (L, l)$ on the unit sphere, we employ standard spherical coordinates, which can be written as $\tilde{\mathbf{s}} = (\cos(l) \sin(L), \sin(l) \sin(L), \cos(L))$. We create a large dataset embedded in the sphere using a target function defined by

$$f(\tilde{\mathbf{s}}, \theta_1, \theta_2, \theta_3) = \theta_1 s_1^3 + \theta_2 s_1 s_2 + \theta_3 s_3^2, \quad (8)$$

where the true values for $(\theta_1, \theta_2, \theta_3)$ are set to $(0.3, 0.7, 0.9)$. The location input (L, l) is generated on the regular grid $1.9^\circ \times 2.5^\circ$ (to simulate over the same grid as in WACCM), which yields a total of $n = 13,824$ location inputs on the sphere. The simulated model output is $\eta(\tilde{\mathbf{s}}, \boldsymbol{\theta}) = f(\tilde{\mathbf{s}}, \boldsymbol{\theta}) + \delta(\tilde{\mathbf{s}})$, where the bias function is $\delta(\tilde{\mathbf{s}}) = 0.05 s_2 s_3$ (around 10% of standard deviation of the model outputs). To investigate the effect of an increase in the number of computer runs on the accuracy of the calibration, we ran $r = 25$ and 50 simulations. The design of the experiments corresponding to the calibration parameters is a maximin LHD. With this design, we tried to cover as much space as possible in the three-dimensional space of the calibration parameters $(\theta_1, \theta_2, \theta_3)$ with 25 and 50 runs, respectively, using the same SH bases to decompose each computer run output, the observational surface and the discrepancy. The SH coefficients are estimated by the least squares method.

The MCMC chain was run for 10,000 samples corresponded for $\boldsymbol{\theta}$, the first 5,000 samples in each chain were discarded as a burn-in and the average computation times (in minutes) for 1 chain were recorded by the Matlab on the Mac OS 10.9.5 platform. The posterior mean and

Table 1: Number of coefficients (N_y) in different orders of the SHT, either for the full expansion (F), or for the restricted expansion (R) assuming $h \geq 0$.

Order	1	2	3	4	5	6	7
F	4	9	16	25	36	49	64
R	3	6	10	15	21	28	36

 Table 2: The posterior mean and SD of θ in function (8) under different computer runs (r) and different orders from (a) full or (b) positive analogue part of SHT with computational times (in minutes).

(a) Full expansion coefficients ($k \geq h \geq -k$)					CPU
r	order	$\theta_1(=0.3)$	$\theta_2(=0.7)$	$\theta_3(=0.9)$	Time
25	1	0.4987(0.2920)	0.6818(0.1276)	0.9004(0.0587)	3
	2	0.2997(0.0092)	0.7919(0.0125)	0.9351(0.0128)	6
	3	0.3133(0.0790)	0.7619(0.0655)	0.9097(0.0335)	11
	4	0.3235(0.1376)	0.7478(0.1121)	0.9193(0.0498)	22
	5	0.3206(0.1555)	0.7232(0.1388)	0.9030(0.0590)	74
50	1	0.5015(0.2887)	0.7021(0.0977)	0.9123(0.0485)	4
	2	0.2893(0.0065)	0.7497(0.0093)	0.9117(0.0096)	13
	3	0.2876(0.0179)	0.7483(0.0246)	0.9039(0.0202)	52
	4	0.2957(0.1050)	0.7249(0.1054)	0.9152(0.0481)	134
	5	0.3066(0.1106)	0.7250(0.1057)	0.9022(0.0498)	388
(b) Reduced expansion coefficients ($h \geq 0$)					CPU
r	order	$\theta_1(=0.3)$	$\theta_2(=0.7)$	$\theta_3(=0.9)$	Time
25	1	0.5043(0.2888)	0.4916(0.2530)	0.8612(0.0819)	2
	2	0.3387(0.0176)	0.7913(0.0228)	0.9372(0.0186)	3
	3	0.3386(0.0150)	0.7892(0.0122)	0.9387(0.0116)	5
	4	0.3266(0.0225)	0.7750(0.0140)	0.9371(0.0132)	11
	5	0.3205(0.1011)	0.7516(0.1095)	0.9021(0.0499)	21
50	1	0.4953(0.2875)	0.5337(0.2386)	0.8651(0.0833)	4
	2	0.2912(0.0077)	0.7483(0.0157)	0.9137(0.0139)	7
	3	0.2912(0.0054)	0.7475(0.0102)	0.9136(0.0097)	16
	4	0.3017(0.0127)	0.7363(0.0117)	0.9131(0.0114)	53
	5	0.2966(0.0360)	0.7381(0.0267)	0.9082(0.0494)	105

standard deviation (SD) for θ calibrated by the 1st to 5th order SH coefficients are presented in the first half of Table 2. Even the 2nd order SHT can adequately capture the variability of the simulation outputs, and the estimation of the true value gets more precise with increasing maximum order of SHs. Nevertheless, the improvement achieved by increasing the order eventually reduces, due to too many coefficients employed to represent θ considering the amount of information available and computation time. The 5th order calibration takes about triple time more than the 4th order calibration in computational cost with similar results. Also, the more computer runs are carried out (50 v. 25), the closer the posterior mean is to the true value under the same order representation. Fig. 3 shows the distribution of posterior means for 100 replicates of an LHD ensemble of size r for calibration parameters (3rd order SHT with $r = 50$), the effect of experimental design on the calibration results is generally small.

Because the normalized SH bases with order k and mode h satisfy $\psi_{k,h}^* = (-1)^h \psi_{k,-h}$, where $*$ denotes complex conjugation, then the same order basis with a different sign in the mode h will have a similar periodic structure. A basis with the same order and a different sign for the mode has the same scale and shape properties, except for the shifting in latitude by $\pi/2$. Therefore we can use only the positive analogue of the SH basis to decompose the model outputs. We can thus reduce the number of SH coefficients from $N_y = (K+1)^2$ to $N_y = (K+1)(K+2)/2$ under the K -th order expansion, see Table 1 for a comparison.

Using the same algorithm and prior specifications, the posterior mean and SD of θ is presented in the second half of Table 2, employing the positive analogue of the SHT. In the lower order expansions, less accuracy is shown than with the complete expansion of the SHs,




Figure 3: Distribution of posterior means for 100 replicates of LHD for the calibration parameters (3rd order SHT with $r = 50$).

due to having less coefficients included. But this approach soon catches up in the higher order expansions. For higher orders, the positive analogue representation outperforms the complete expansion of SHs at a much lower computational cost. Indeed, from a statistical point of view, this approach avoids fitting redundant information, and aids identification.

For the estimation of model discrepancies, we use a full expansion of SHs with $r = 50$ as illustrations. We also employ a varying number of basis representation for y^F and η to investigate the impact on the reconstruction of model discrepancies as well as on the estimation of calibration parameters. Fig. 4(a) displays the discrepancies we added to the model outputs, and Fig. 4(b) shows the estimated mean discrepancies reconstructed from the 2nd order SHT for y^F and η . The large scale behavior is captured, but there is misfit as well, in particular at both poles. For a 3rd expansion order to y^F , combined with 3rd or 2nd order to η , Fig. 4(c) and (d) both provide a good reconstruction of model discrepancies.

On the other hand, employing different levels of expansion in the basis functions for y^F and η is not necessary to improve the estimation of calibration parameters. The results of posterior means and SDs for θ are presented in the Table 3. From this table we can see that the posterior means estimated from higher expansion order become numerically unstable. Indeed, the model outputs are sensitive to the higher order representations, and this information is truncated from the covariance structure. Model D–E and G–H show a similar pattern in underestimation of θ_1 slightly and θ_2 significantly, which means that certain bases between 3rd to 4th are important for quantifying the variability of the simulation outputs, and the variability was included in Model I so as help the estimation. Computational costs increase in the higher order representations, due to low dimensional matrix that cannot be accelerated by the sparse matrix manipulation. In real applications we should carefully transform the spatial output on the balance of the estimation of the calibration parameters and flexibility on the reconstruction of model discrepancy.

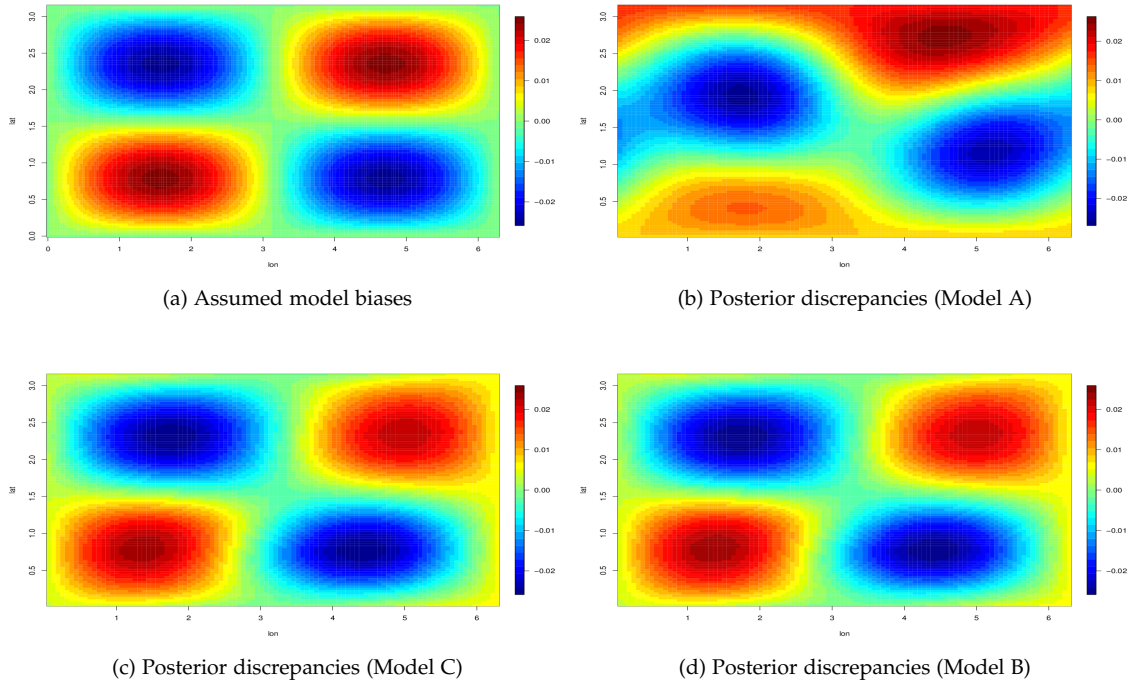
3.3 Example 3: Comparison with principal components

We now compare the performance of our approach to that of the PC approach proposed by Higdon et al. (2008). The function with three calibration parameters is

$$f(\tilde{s}, \theta) = \theta_1 s_1 + \exp(-\theta_2 s_2 - \theta_3 s_3^2), \quad (9)$$

Table 3: The posterior mean and SD of θ in function (8) from SHT with varying expansion orders for model outputs and observations with computational times (in minutes).

Model	order		$\theta_1(=0.3)$	$\theta_2(=0.7)$	$\theta_3(=0.9)$	CPU Time
	η	y^F				
A	2	2	0.2893(0.0065)	0.7497(0.0093)	0.9117(0.0096)	13
B	2	3	0.3056(0.0289)	0.7305(0.0344)	0.9016(0.0254)	50
C	3	3	0.2876(0.0179)	0.7483(0.0246)	0.9039(0.0202)	52
D	2	4	0.2401(0.0942)	0.5819(0.1014)	0.9106(0.0509)	131
E	3	4	0.2322(0.0935)	0.5813(0.1026)	0.9153(0.0506)	130
F	4	4	0.2957(0.1050)	0.7249(0.1054)	0.9152(0.0481)	134
G	2	5	0.2672(0.1028)	0.6072(0.1039)	0.9162(0.0479)	309
H	3	5	0.2837(0.1434)	0.6073(0.1456)	0.8884(0.0651)	353
I	4	5	0.3081(0.0959)	0.7364(0.0974)	0.9113(0.0452)	338
J	5	5	0.3066(0.1106)	0.7250(0.1057)	0.9022(0.0498)	388


 Figure 4: (a) Assumed biases surface; posterior mean model discrepancies from the 2nd (Model A) and 3rd (Model C) order SHT for both η and y^F , and from the 2nd and 3rd order SHT for η and y^F , respectively (Model B).

where the true value for $(\theta_1, \theta_2, \theta_3)$ is set to $(0.3, 0.7, 0.9)$. We used locations $(L_i, l_i), i = 1, \dots, 100$ on a regularly spaced grid. The observational surface is $y^F(\tilde{s}, \theta^*) = f(\tilde{s}, \theta^*) + \epsilon(\tilde{s})$, where $\epsilon \sim N(0, 0.01)$, and $r = 50$ simulator runs are obtained according to a maximin LHD, denoted by $\theta_j = (\theta_{1j}, \theta_{2j}, \theta_{3j}), j = 1, \dots, 50$.

The PC based calibration algorithm developed by Higdon et al. (2008) performs a dimensional reduction across the output space. This projection is constructed by a data-driven framework (singular value decomposition). The SHT provides a different approach: it reduces the size of each model output using the same basis functions to decompose each computer run's output into coefficients. The decomposition is also applied to the observations in order to compare it with the output in the SHs domain.

For the PCs approach, we cannot directly apply the algorithm from Higdon et al. (2008) due to the lack of SH basis functions for the joint input-output projection. Therefore we modify their framework to compare our method with the essential part of their approach. The output surfaces η are decomposed by the empirical orthogonal functions (EOFs). EOFs are merely geographically weighted principal components. The EOFs are found by computing the eigenvalues and eigenvectors of a spatially weighted anomaly covariance matrix of a field. The derived eigenvalues provide a measure of the percentage of variance explained by each mode. We calculate the singular value decomposition (SVD) for the model outputs:

$$\begin{matrix} & \theta_1 & \dots & \theta_{50} \\ s_1 & \left(\begin{array}{ccc} \eta_{1,1} & \dots & \eta_{1,50} \\ \vdots & \ddots & \vdots \\ \eta_{100,1} & \dots & \eta_{100,50} \end{array} \right) \\ \vdots & & & \\ s_{100} & & & \end{matrix} = UDV^T,$$

where the columns of U are the left singular vectors; D has singular values and is diagonal; and V^T has rows that are the right singular vectors. The SVD represents an expansion of the original data in a coordinate system where the covariance matrix is diagonal. The EOFs are then the leading p_η columns of $\frac{1}{\sqrt{r}}UD$.

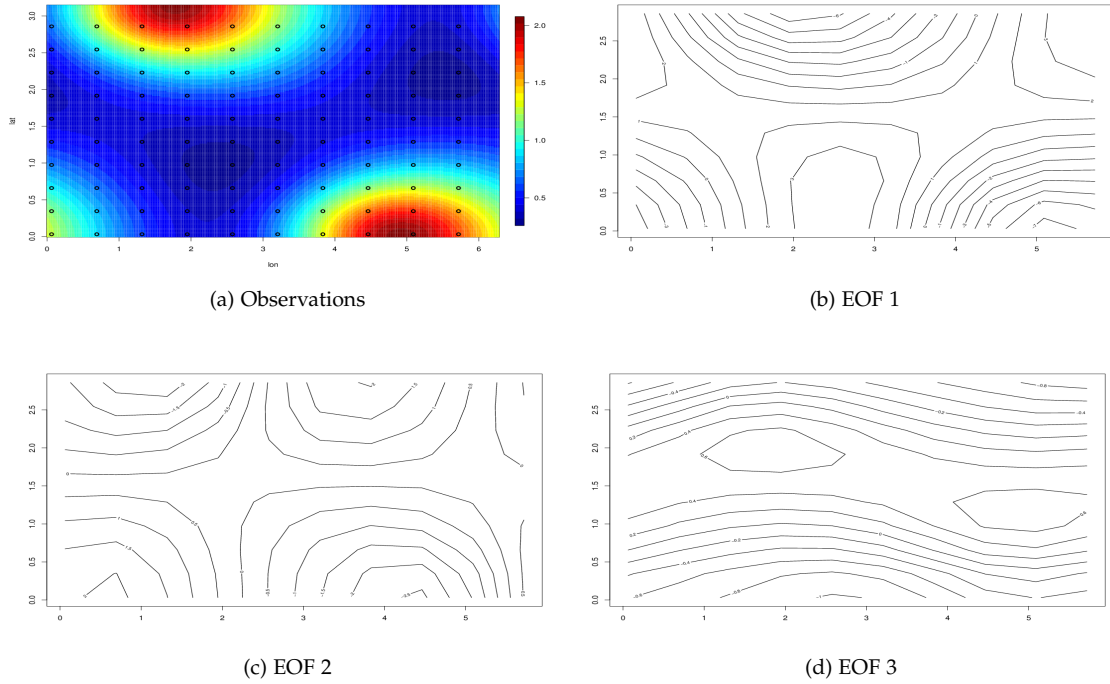


Figure 5: Observations with measurement locations (circles) and first three EOFs.

We retain the first $p_\eta = 10$ EOFs: the observations and contour of the first three EOFs for the model outputs are presented in Fig. 5. Indeed, the cumulative variations explained by the first three components are 71.28%, 82.82% and 92.76%. For the calibration, we also assume that EOF weights are uncorrelated. The GP assumption was imposed on each EOF weight $w_i^p(\mathbf{s}, \boldsymbol{\theta})$

with mean 0 and covariance structure


$$\text{Cov}((i, \theta), (i', \theta')) = \frac{1}{\lambda_{wi}} \prod_{k=1}^3 \rho_{wi(k+2)}^{4(\theta_k - \theta'_k)^2} \times I_{ii'}. \quad (10)$$

The discrepancy function is assumed to be normal with mean 0 and bias precision λ_δ . The assumptions for priors are the same as the SH approach: $\lambda_\epsilon \sim \Gamma(1, 0.003)$, $\lambda_\delta \sim \Gamma(1, 0.01)$, $\lambda_{wi} \sim \Gamma(5, 5)$, $i = 1, \dots, p_\eta$, $\rho_{wk} \sim \text{beta}(1, 0.1)$, $k = 1, \dots, 3$ and uniform priors for θ . As for the SHs approach, we performed the first five order transforms for the observation and each computer run (see Fig. 6). To illustrate the difference with the EOF approach, we use the 3rd order expansion (16 coefficients) as an example. Our approach represents the input–output relationship as follows


$$\begin{array}{ccc} \text{input} & & \text{output} \\ \begin{pmatrix} L_1 & l_1 & \theta_{1,1} & \theta_{1,2} & \theta_{1,3} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ L_m & l_m & \theta_{m,1} & \theta_{m,2} & \theta_{m,3} \end{pmatrix}_{(m \times (2+3))} & \xrightarrow{m=n \times r} & \begin{pmatrix} \eta_1 \\ \vdots \\ \eta_m \end{pmatrix}_{(m \times 1)} \\ \text{input} & & \text{output} \\ \begin{pmatrix} z_1 & \theta_{1,1} & \theta_{1,2} & \theta_{1,3} \\ \vdots & \ddots & \vdots & \vdots \\ z_{m'} & \theta_{m',1} & \theta_{m',2} & \theta_{m',3} \end{pmatrix}_{(m' \times (1+3))} & \xrightarrow[n \gg N_\eta]{m' = N_\eta \times r} & \begin{pmatrix} c_1^M \\ \vdots \\ c_{m'}^M \end{pmatrix}_{(m' \times 1)} \end{array}$$

basis representations $\{\psi_z\}$

where $m = 5000$, $n = 100$, $r = 50$, $m' = 800$ and $N_\eta = 16$. Our approach not only reduces the output dimension, but also changes the input settings (from spatial inputs to basis function indices). The model output η is transformed into the coefficients $c_z^M(\theta)$, with index z as an univariate input and θ as continuous calibration inputs. The coefficients c_z^M obey the GP assumption with mean 0 and covariance structure defined in Equation (1).



(a) Design on the S^2



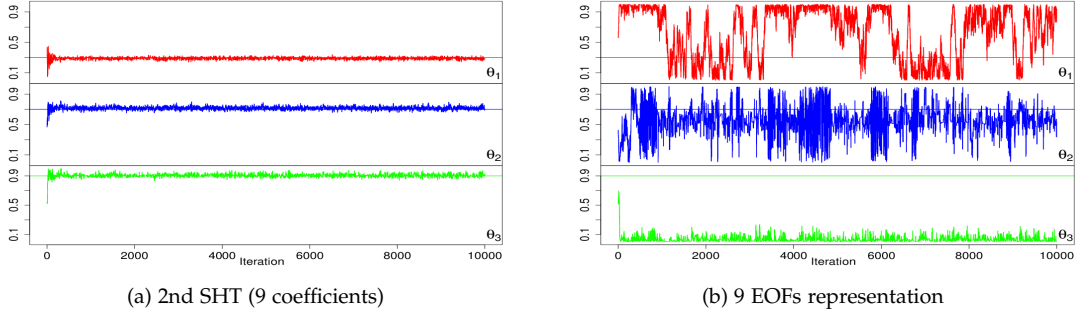
(b) Estimated 3rd SH coef.

Figure 6: Observations (circles) and model output (crosses) on the sphere (a) and estimated 3rd order SH coefficients (b).

The posterior mean and SD for the calibration parameters are displayed in Table 4. The best estimation of θ^* in the SHs approach is with the 3rd order representation. The posterior densities approximated by the 4th order representation present higher uncertainties. The EOF approach fails to converge towards the true values for the cases of 1 to 10 components. Fig. 7 displays the MCMC sample path of the calibration parameters θ for the 2nd order SHs (9

Table 4: The posterior mean and SD of θ for the empirical orthogonal functions (EOFs) and spherical harmonics (SHs) approaches.

	components/orders	$\theta_1(=0.3)$	$\theta_2(=0.7)$	$\theta_3(=0.9)$
EOF	1	0.4955(0.2928)	0.5014(0.2905)	0.5088(0.2921)
	2	0.5025(0.2957)	0.4914(0.2909)	0.2957(0.2861)
	3	0.4964(0.2944)	0.5095(0.2811)	0.4983(0.2892)
	4	0.5851(0.3591)	0.4957(0.2212)	0.0274(0.0289)
	5	0.6085(0.3917)	0.6030(0.2754)	0.0612(0.0764)
	6	0.3916(0.3131)	0.7610(0.1462)	0.0600(0.0603)
	7	0.3651(0.3273)	0.6708(0.2099)	0.0427(0.0464)
	8	0.4199(0.3206)	0.6894(0.1644)	0.0341(0.0419)
	9	0.4441(0.3400)	0.5967(0.2020)	0.0387(0.0412)
	10	0.4361(0.3684)	0.5745(0.2312)	0.0345(0.0450)
SH	1	0.5180(0.2826)	0.7113(0.1237)	0.9191(0.0526)
	2	0.2899(0.0055)	0.7156(0.0079)	0.9094(0.0083)
	3	0.2938(0.0045)	0.7035(0.0570)	0.9037(0.0336)
	4	0.2911(0.1026)	0.6880(0.1008)	0.9081(0.0514)
	5	0.3250(0.1155)	0.7169(0.1094)	0.9012(0.0509)


 Figure 7: MCMC paths of θ for SHs and EOFs algorithms. Solid lines indicate the true values.

coefficients) and 9 EOFs representation. We can see that for all calibration parameters in the SHs approach, convergence occurred after roughly 500 iterations, and convergence of the chains cannot be established for the EOF approach.

The choice of a suitable basis representation is usually highly application dependent. Note that SHs and B-spline are location-dependent basis functions, while EOF bases are data driven. The limitation of our technique is that SHT can only be applied to spatial data, whereas EOF decomposition are more widely applicable (but do not acknowledge the nature of the spaces on which it is performed: EOFs do not explicitly account for spatial correlation, only vectorize the spatial outputs). Therefore we can not conclude that the SH basis approach is always a better choice.

3.4 Example 4: Nonstationary field

We now illustrate how the parameters in an SPDE approach can be incorporated into our calibration algorithm to model nonstationarity over a spherical domain. With 10×10 regularly spaced locations in latitude and longitude and 50 computer runs according to a maximin LHD for the calibration inputs, the function with three calibration parameters is set to

$$f(\tilde{\mathbf{s}}, \boldsymbol{\theta}) = \left(\frac{1}{2} s_1^2 + \theta_1 s_2 s_3 \right) \times \begin{cases} \theta_2 s_2 & \text{if } L \leq \pi/2 \\ \theta_3 \exp(-s_3 - s_1) & \text{if } L > \pi/2 \end{cases}, \quad (11)$$

where the true values for $(\theta_1, \theta_2, \theta_3)$ are set to $(0.5, 0.2, 0.8)$. We create a nonstationary spatial field by introducing different structures in the Northern and Southern hemispheres, where θ_1 is a global calibration parameter, and (θ_2, θ_3) are local variates, see Fig. 8. First, we perform the SHT onto observations y^F and each computer run $\eta_j, j = 1, \dots, 50$, and then carry out the calibration on the coefficients.




Figure 8: Observations for function (11).

Here we only use the positive analogue $h > 0$ basis representation. In total, we estimate 13 models with different numbers of expansion order. The results of the analysis of the calibration parameters using the 4th to 7th orders of the SHT are shown in the first part of Table 5 (Model A-D). From the results for the posterior sampling, we can see that the global calibration parameter θ_1 is estimated well. However, even though the convergence of an MCMC chain can be established for θ_2 and θ_3 , the posterior means are underestimated. Besides, an increase in the order of the expansion does not really improve the results. This underestimation can be viewed as an inability to capture local variations. As $s_2 = \sin(l) \sin(L)$ has stronger variations than $\exp(-s_3 - s_1) = \exp(-\cos(L) - \cos(l) \sin(L))$ (see Fig. 8), those numerical differences also influence the estimation of θ_2 and θ_3 .

Table 5: Posterior mean and SD for $(\theta_1, \theta_2, \theta_3)$ in function (11), and number of coefficients (right column) under different orders of SHT for $\{\eta, \kappa, \tau\}$ per model run.

Model	η	κ	τ	$\theta_1(=0.5)$	$\theta_2(=0.2)$	$\theta_3(=0.8)$	$N_\eta + N_\kappa + N_\tau$
A	4	-	-	0.5054(0.0503)	0.1884(0.0479)	0.7620(0.0384)	15
B	5	-	-	0.4975(0.0530)	0.1794(0.0622)	0.7462(0.0499)	21
C	6	-	-	0.4765(0.0619)	0.1664(0.0791)	0.7050(0.0685)	28
D	7	-	-	0.4880(0.1121)	0.1976(0.1265)	0.6954(0.1187)	36
E	-	1	1	0.5793(0.1578)	0.1475(0.0684)	0.6201(0.1997)	6
F	-	2	2	0.5603(0.0969)	0.1887(0.0782)	0.7396(0.0892)	12
G	-	3	3	0.7852(0.0782)	0.4423(0.1554)	0.8581(0.0536)	20
H	4	1	1	0.4516(0.0970)	0.0708(0.0489)	0.4948(0.0367)	21
I	5	1	1	0.4950(0.0438)	0.1331(0.0489)	0.4984(0.0321)	27
J	6	1	1	0.3558(0.0500)	0.1347(0.0521)	0.6855(0.1193)	34
K	4	2	2	0.5532(0.0680)	0.2245(0.1083)	0.7711(0.1094)	27
L	5	2	2	0.5287(0.0676)	0.1785(0.1067)	0.7935(0.0978)	33
M	6	2	2	0.5370(0.0659)	0.1713(0.1098)	0.7894(0.0828)	40

In order to understand the role of the parameters of the SPDE in the calibration, we first perform a calibration using only the SPDE estimated coefficients $\{\kappa^M, \tau^M\}$. Under the same priors and algorithm, the posterior mean and SD of the first three orders of the expansion for κ^M and τ^M are shown in the second part of Table 5 (Model E - G). Even though the calibration does not fully succeed (and should not without matching original outputs to observations but

only SPDE information), the result in the 2nd order expansion for κ^M and τ^M seems informative as the posterior modes are close to the true values. The first two orders of the expansion surface for κ^F and τ^F for the observations y^F are shown in Fig. 9. It is difficult to directly interpret the features of $\kappa(s)$ and $\tau(s)$; However, from Fig. 9(c) - (d) we can see a strong northeast-southwest flow with high uncertainty (inverse precision) and high correction decay rate over the y^F surface. Besides, the 1st and 3rd order representation results seem unhelpful in detecting the true value of the calibration parameters.

For the next step, we infer $\{c^M, \kappa^M, \tau^M\}$ jointly with the GP model. We assumed that these three types of coefficients are uncorrelated. Therefore, the covariance matrix in the GP model remains block diagonal. The results of the analyses for the 4th to 6th order representations for the observations and model outputs with the first two orders of expansion with the parameters in the SPDE included are presented in the third part of Table 5. We can see that with the parameters of the SPDE included, we achieve an improvement in the calibration. For example, Model C and K have a similar number of coefficients, but the combination with SPDE increases the estimated accuracy in θ_2 and θ_3 . The similar case of Model D and L also supports the use of SPDE information. Model L uses a smaller number of coefficients, while achieving an improvement in terms of increased accuracy in θ_3 and reduced posterior uncertainties. Nevertheless, only in the case of the 2nd order expansion do the SPDE parameters help; the 1st order expansion cannot achieve a good result.

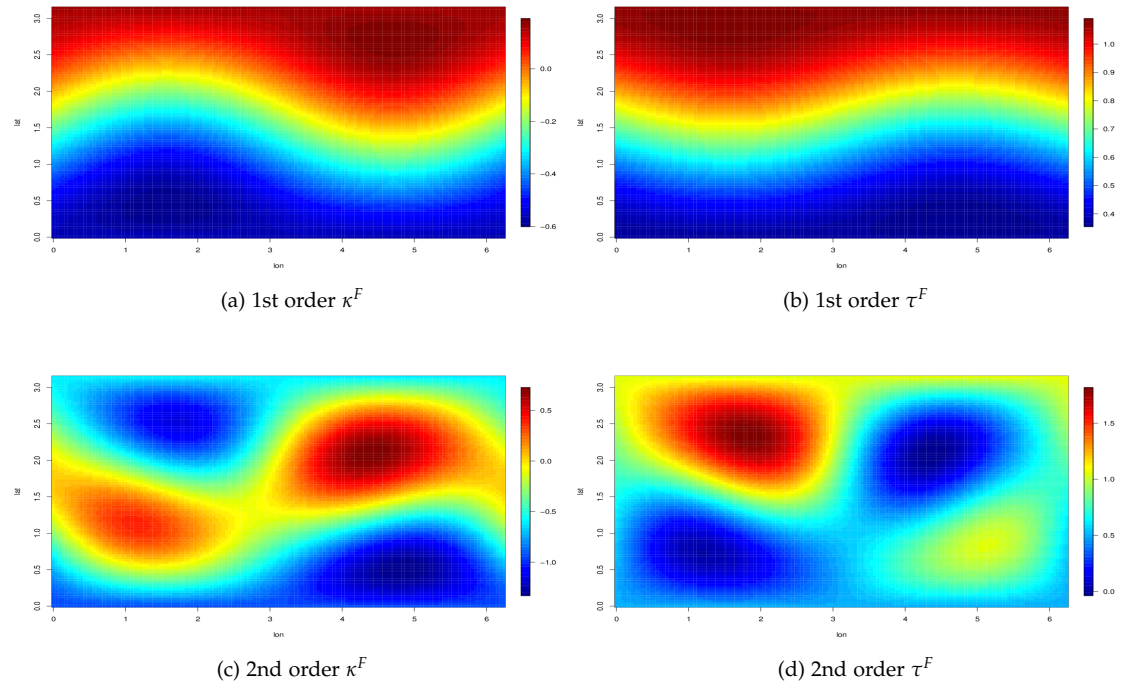


Figure 9: The first 2 order expansion surface for κ^F and τ^F for observations y^F .

The calibration with the SPDE defined scale and precision parameters did achieve an improvement in terms of increasing the accuracy of the MCMC sampling. However, the impact of the SPDE information has been inconsistent. It also augments uncertainties in the posterior densities in some cases. It is necessary to achieve a compromise between accuracy and precision. Besides, in real applications, we often do not know whether the calibrated values work until actually performing a validation. It can be computationally challenging to find the optimized

orders for the combination of η , κ and τ .

Similar to most truncated basis representations, we choose the number of basis functions *post hoc*. Our experience can provide model selection criteria as follows: (1) A basis representation for model outputs η play a dominant role in the algorithm. The number of basis functions for the model outputs (N_η) usually needs to be greater than the total number of basis functions for scale and precision parameters ($N_\kappa + N_\tau$); (2) Calibration with only one of the coefficients κ or τ cannot improve the analysis. The reason is unclear, but the fact that κ and τ represent a spatial process jointly is tacitly assumed. Recall that the Matérn function is controlled by the smoothness parameter ν , the scale parameter κ , and the precision parameter τ . The smoothness parameter ν is fixed by $\alpha = \nu + d/2$ in connection with the SPDE, therefore the approximated spatial process depends upon κ and τ jointly. Both κ and τ need to be included to reflect the full variation in the spatial field.

3.5 Application to the WACCM experiments

A series of WACCM runs with the component set prescribed sea ice, data ocean and specified chemistry, with horizontal resolution $1.9 \times 2.5^\circ$ and 66 vertical levels were simulated from 1st January 2000. The GW parameterizations in WACCM depend on four inputs: (1) *cbias* ($\theta_1 \in [-5, 5]$): anisotropy of the source spectrum, e.g. -5m/s: the spectrum has a stronger westward component, with center of the spectrum at 5m/s westward; (2) *effgw* ($\theta_2 \in [0.05, 0.3]$): the efficiency factor, measures the gravity wave intermittency; (3) *flatgw* ($\theta_3 \in [1, 3]$): controls the momentum flux of the parameterized waves at the launch levels; (4) *launlvl* ($\theta_4 \in [50, 700]$): launch levels of the waves. The value of GW inputs θ are generated by a maximin LHD. We simulated $r = 100$ runs for 2 months. The first month was discarded as a spin-up period. Each output was computed over 96 latitudes and 144 longitudes, so the dimension is $n = 96 \times 144 = 13,824$.

3.5.1 Calibration against synthetic observations

To illustrate our methodology, we compare the zonal wind simulations $\eta(\mathbf{s}_i, \theta_j)$, where $\mathbf{s}_i, i = 1, \dots, 96 \times 144$ are the latitude and longitude on the spherical domain, and $j = 1, \dots, 100$ is the index of the runs, with outputs from WACCM's standard case, instead of actual observations. Therefore we know the true GW parameters values and can validate our method. Let $\eta^*(\mathbf{s}_i)$ be the zonal wind surface from WACCM standard output. In order to account for possible observation error and lack of physics in the model (discrepancy), we add a smooth noise to $\eta^*(\mathbf{s}_i)$ by assuming that the observations are

$$y^F(\tilde{\mathbf{s}}_i) = \eta^*(\tilde{\mathbf{s}}_i) + \frac{\sigma_{\eta^*}}{5}s_1 + \frac{1}{2}s_2s_3,$$

where $\tilde{\mathbf{s}}_i = (s_1, s_2, s_3)$ are the spherical coordinates, and $\sigma_{\eta^*} = 11.14$ is the SD of η^* . Fig. 10(a) and (b) shows the zonal wind surfaces from standard outputs and synthetic observations at 30mb, February 2000.

For ease of implementation of the MCMC algorithm, we standardize the entire set of responses (observations and WACCM outputs) by the mean and SD of the WACCM zonal wind outputs before performing the SH decomposition. Therefore the coefficients can be assumed to be mean 0 without loss of generality and the marginal variance in the model outputs is approximately 1. The design space of the calibration parameters is also scaled to be $[0, 1]^4$. As for the computational issue, in practice it is difficult to deal with a size of model output beyond moderately large (say $\simeq 2000$ responses). Here we have $r = 100$ computer runs, therefore we seek to decompose each model output with about 20 coefficients. To allow more flexibility for representing observations and model discrepancies, the results are illustrated by the 3rd and 4th

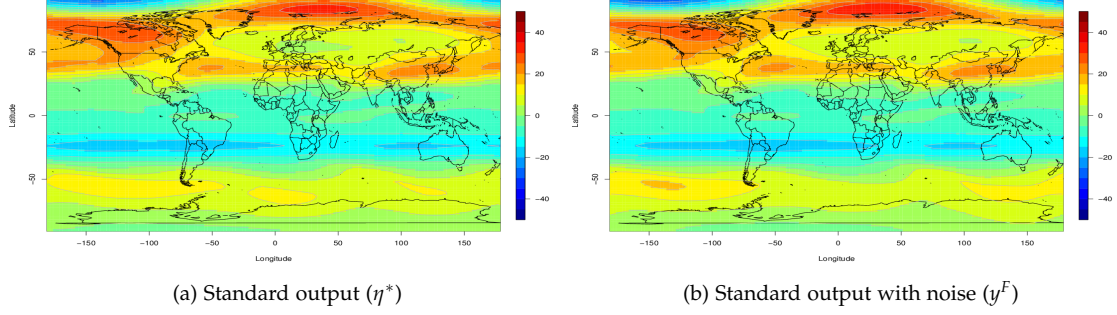


Figure 10: (a) Zonal wind standard output; (b) Assumed observed surface: noise and discrepancy added to the zonal wind standard output (30mb, Feb. 2000).

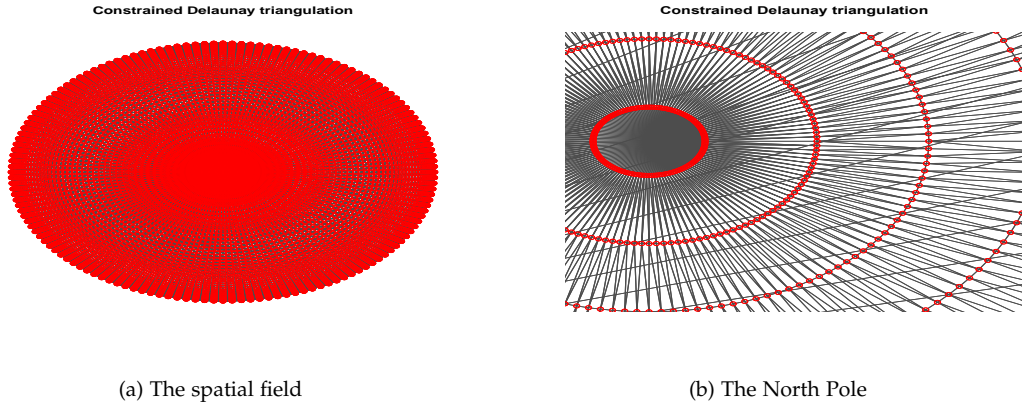


Figure 11: Triangulation of the spatial field. Red points indicate the output locations.

order decomposition in SH bases for model outputs and observations, respectively, including 1st order SPDE defined precision and scale parameters. As the WACCM simulator produces output on a high resolution regularly spaced grid over a sphere, we construct a spherical triangulation by the INLA–SPDE default algorithm without any other particular considerations, and both model outputs and observations share the same triangulation. The triangulation of the WACCM output domain is provided in Fig. 11.

All the unknown parameters in the algorithm require specified prior distributions which represent uncertainty about the values of these parameters before any data is collected. The following choices are made for the priors: (a) We specify an uniform prior distribution over each of the GW parameter interval; (b) To model the correlation parameters $\rho_{\eta_k}, k = 1, \dots, 4$, a Beta(1, 0.1) distribution is used, which conservatively places most of its prior mass on values of ρ_{η} near 1 (indicating an insignificant effect); (c) Gamma prior distributions are used for each of the precision parameters $\lambda_{\eta}, \lambda_{\delta}$ and λ_{ϵ} . Specifically, we use priors $\lambda_{\eta} \sim \text{GAM}(5, 5)$ (with expectation 1 due to standardization of the responses), $\lambda_{\delta} \sim \text{GAM}(1, 0.01)$ (with expectation around 10% of SD of the standardized responses) and $\lambda_{\epsilon} \sim \text{GAM}(1, 0.003)$ (with expectation around 5% of SD of the standardized responses). We implement a Metropolis–Hastings algorithm to explore the multidimensional space of parameters. Metropolis updates are used for the correlation and the calibration parameters with a uniform proposal distribution centered at the current value of

the parameter. The precision parameters are sampled using Hastings updates with a uniform proposal distribution centered at the current value of the parameter (Higdon et al., 2008). This eventually yields draws from the posterior distribution by repeatedly accepting and rejecting a choice of move in the parameter space.

Fig. 12(a) shows the boxplots of the marginal posterior distributions for the ρ_η 's, which control the dependence strength in each pair of θ in the GP model. The posterior density of ρ_4 closes to 1 that indicate a very weakly significant effect for θ_4 . The marginal posterior densities of each calibration parameter are displayed in Fig. 12(b). The posterior modes are (0.4347, 0.5473, 0.0597, 0.2762) for each of the calibration parameter in $[0, 1]$ scale. True values of default GW parameters are (0.5, 0.56, 0, 0.2308). Our technique shows the ability to estimate θ_3 , a difficult task as the true value lies on the lower bound. Our approach provides a good compromise between computational feasibility and fidelity to the data by only using parsimonious spherical representations. The results suggest that our technique on calibration of global-scale outputs is effective.

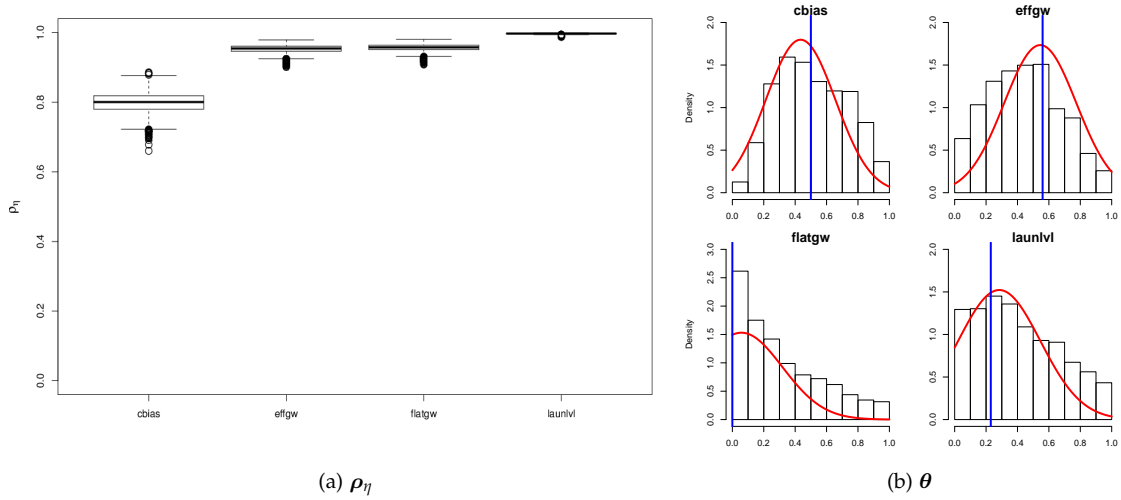


Figure 12: (a) Boxplots of the marginal posterior distribution for correlation parameters ρ_η ; (b) Marginals for the posterior distribution of the GW parameters θ . Vertical lines indicate the true values (30mb, Feb. 2000).

3.5.2 Calibration against real observations

The final step is to carry out the calibration against real observations. We use zonal wind data obtained from the *European Centre for Medium-Range Weather Forecasts* (ECMWF) 40 Year Re-analysis (ERA) Data Archive. We focus on the altitude of 1mb, as the outputs in low altitudes are less sensitive to GW parametrizations and match the observations already well. Fig. 13(a) and (b) shows the ERA observations and zonal wind surfaces from standard outputs at 1mb, February 2000. Fig. 13(c) shows the zonal means calculated over every 5° belt of observations (black solid line), standard outputs (black dashed line) and each run of model output (grey dotted lines): note that the zonal mean winds over the Tropics are high compared to the observations and standard outputs. The input value of θ_3 in the standard output is being at the lower border of parameter range, this may produce relatively extreme behavior over the Tropics in our model runs. Fig. 13(d) represents the grid-by-grid SDs map across model outputs. We can see that the spatial process is clearly anisotropic and highly latitude dependent; the

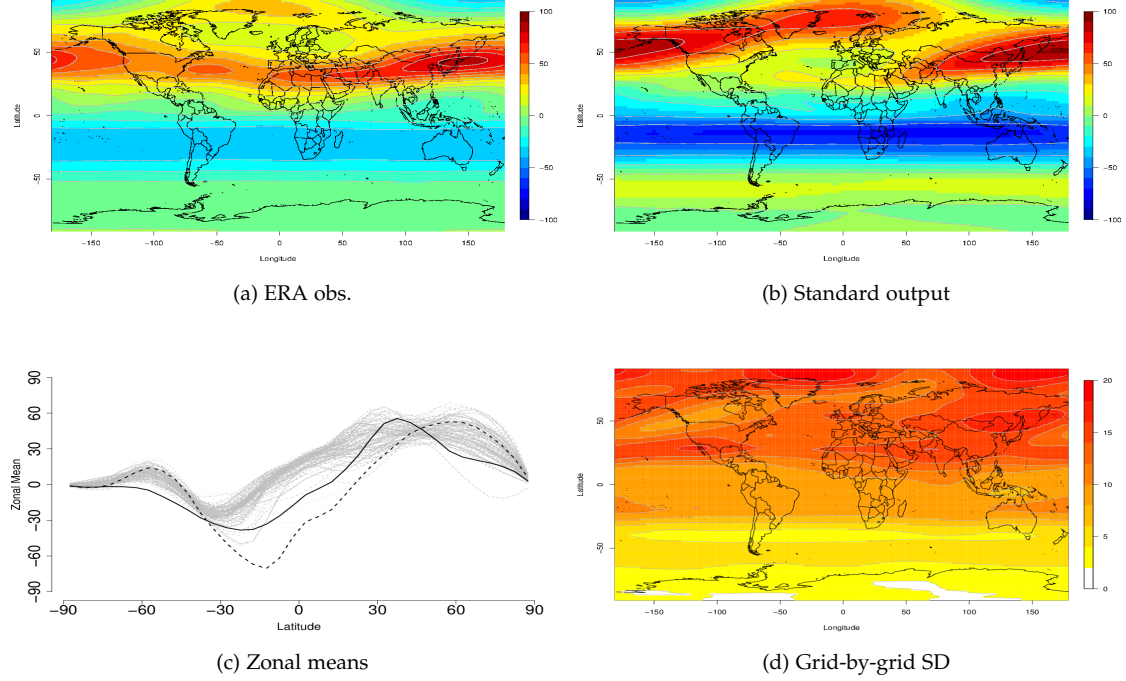


Figure 13: Zonal wind from the (a) ERA data and (b) WACCM standard output; (c) Zonal means of observations (black solid line), standard (black dashed line) and model outputs (grey dotted lines); (d) Grid-by-Grid SDs map across model runs (1mb, Feb. 2000).

uncertainties are concentrated over the Northern Hemisphere, and little significant variabilities can be found over the Southern Hemisphere.

Under the same model settings described in the previous section, Fig. 14 (a) shows the sample paths for 3 chains, with 6000 iterations, corresponding respectively to the calibration parameters. The convergence of the MCMC chain can be established for the parameters θ_2 , θ_3 and θ_4 , with posterior modes 0.1067 (SD=0.0514), 0.0805 (SD=0.0286) and 0.3389 (SD=0.0175) in the $[0, 1]$ scale, respectively. Besides, the sample paths of θ_1 are hitting the upper boundary constantly. It means that the possible calibrated value may lie outside the boundary. We then use posterior modes for these paths, collected as input values for the validation of WACCM. The zonal wind output surface displayed in Fig. 14(d), shown a root-mean-square error (RMSE) of 18.15, which is a percentage of improvement of 14.99% over the standard output (the RMSE between ERA observations and standard output is 21.35).

Fig. 14(c) compares the differences between observations and mean structure of model outputs in each cell (with respect to 100 LHD.), i.e. $\delta_{\text{initial}}(\mathbf{s}) = y^F(\mathbf{s}) - \bar{\eta}(\mathbf{s}, \theta)$, where $\bar{\eta}(\mathbf{s})$ is the output means over space. This figure provides potential features of model discrepancy over space (albeit not the true discrepancy). As expected, the model tends to overestimate the values over the Tropics, which matches the pattern in Fig. 13(c). Besides, this surface seems to match the pattern in Fig. 13(d). The largest model bias (apart from the Tropics) and variability both occur over Northeast Asia and the North Pole. Fig. 14(d) shows the posterior mean discrepancies surface in the sense of $\delta^*(\mathbf{s}) = y^R(\mathbf{s}) - \eta(\mathbf{s}, \theta^*)$. Our calibration reduces the bias (i.e. overestimation) over the Tropics, as well as the bias (i.e. underestimation) over the North Pole, whereas the bias over the Northeast Asia remains.

We further use posterior modes to simulate 5 years (2 quasi-biennial oscillation cycles)

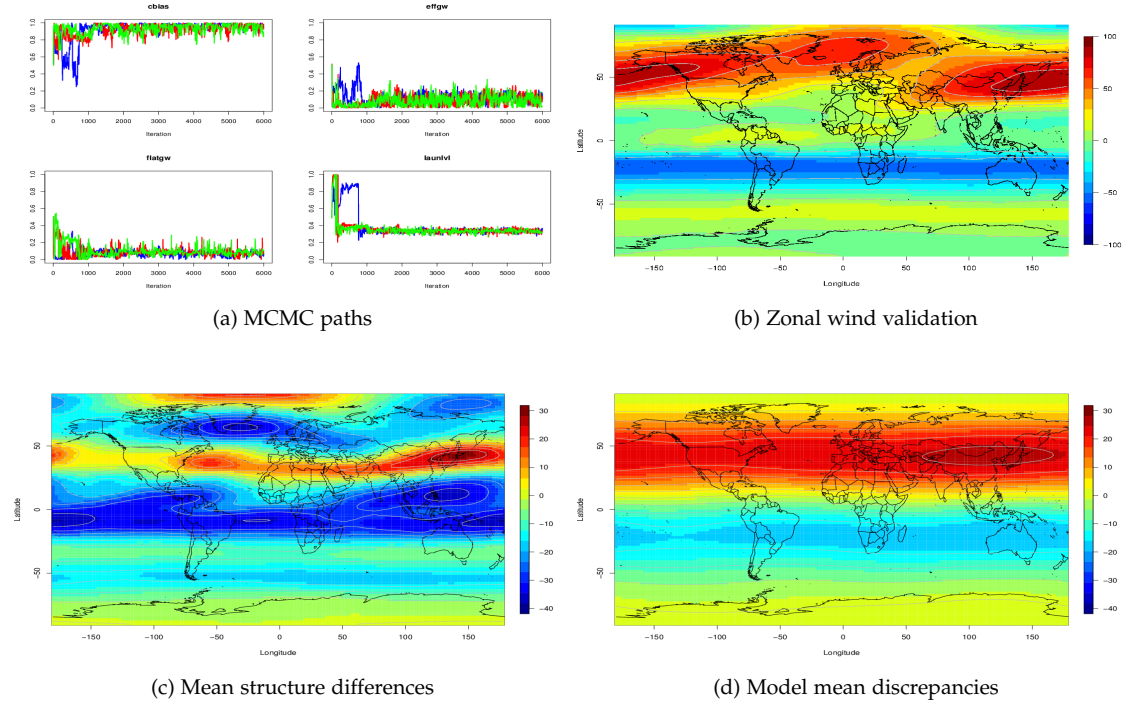


Figure 14: (a) MCMC paths for 3 chains and (b) Zonal wind map generated by posterior from the calibration; (c) Differences between observations and mean structure of model outputs and (d) Model mean discrepancies map (1mb, Feb. 2000).

of zonal wind data. Fig. 15(a) shows monthly RMSEs at 1mb globally, from 2000 to 2004. The overall averaged RMSE for the standard and the calibrated outputs are 24.51 and 22.99, respectively. As the GW parametrization affects the zonal wind over the Tropics more than on the global average, we also investigate RMSEs over the Tropics over the same period. The RMSE trends are shown in Fig. 15(b). The overall averaged RMSE over the Tropic for the standard and the calibrated outputs are 26.64 and 17.87, respectively. Therefore the improvement is more significant over the Tropics, with percentage of improvement 32.9%. Simulation by our calibrated outputs outperform the standard code in 51 months out of 60 months. The results show that our technique is robust to a specific horizontal output from the WACCM simulator. The calibration of WACCM with real observations over the whole output domain constitutes another level of complexity that needs joint scientific and statistical expertise. It is currently under investigation, but is beyond the scope of this paper. Indeed, observations are scarce at these altitudes and show features that require specific understanding of the upper atmosphere dynamics before being used for calibration, and over many years of simulation for an adequate comparison.

4 Conclusion and discussion

Our approach improved the calibration of large-scale computer model outputs distributed over spaces, parsimoniously, by using bases representations of the surfaces over the plane or the sphere. In addition, the INLA-SPDE approach was used to decompose its own parameters over the the same bases in order to improve calibration. The synthetic and real examples confirm

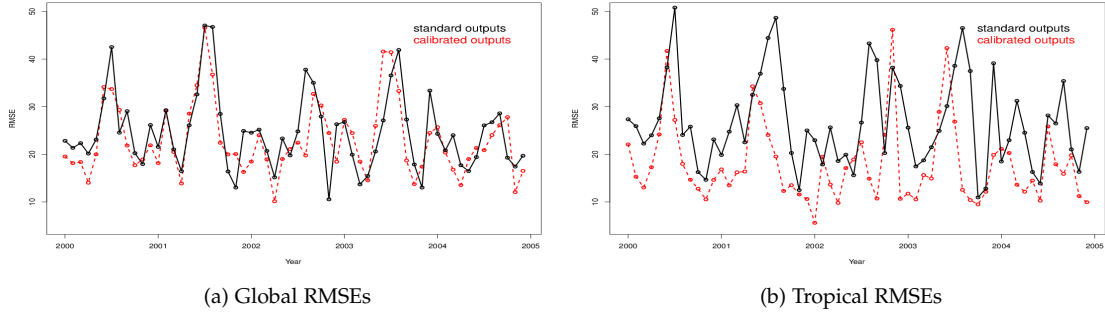


Figure 15: Monthly RMSE trends between the ERA observations and standard outputs (solid line) or calibrated outputs (dashed line), from 2000 to 2004.

the ability of our approach to efficiently and accurately perform calibration. Our method was inspired by the wavelets method of Bayarri et al. (2007), but with a different type of outputs: spatial v. time series. We can expect that the spherical wavelet decomposition may also be a possible alternative basis representation on the spatial domain, whenever appropriate (e.g. sharp variations). We also considered data-driven bases representations, such as principal components analysis, which need to seek enough variation to perform an effective analysis; our technique outperforms principal component analysis in some settings.

On the sphere, the SH coefficients represent the wave features at different scales. For the purpose of calibration, it is unnecessary to approximate the spatial processes with very high order series expansions of SHs to fit each run of model output best. The critical requirement is to extract sufficient variations behind the SH coefficients about the calibration parameters under a suitable basis representation. Muir and Tkalčić (2015) utilized the corrected Akaike information criterion (AIC) to choose an optimal maximum order of expansion for an irregular data on the sphere in a hierarchical Bayesian setting. The results show the 3rd to 5th order of expansion in SHs are generally a turning point from fast to slow reduction in AIC in terms of balancing explanatory power with simplicity (although not the smallest AIC). In all these approaches, the choice of the number of basis vectors is currently *post hoc*. We reckon the 3rd order of expansion in SHs for the model outputs as a good start in practical application.

However, one may miss important small scale features with a low order expansion. In case the higher order SHs are required to capture small but meaningful features in a spatial process, one possible solution is to restrict the mode h to a low order in the series expansion of SHs in order to allow more flexibility in variations across latitudes than across longitudes within latitudes (Stein, 2007). The other possibility is to perform a variable selection to those SH coefficients before calibration. For example, in Figure 6(b), it can be seen that variations present only in certain bases. In this case we can expect that the calibration only through the 1st and 5th to 8th coefficients should be able to provide a similar result (all the coefficients decay after the 9th coefficient). Yet a reliable variable selection technique is required to identify the threshold for the cut-off of redundant coefficients.

Another advantage of using the SH basis, compared to data-driven ones such as PCs, is that sequential design is allowed (Beck and Guillas, 2016), because the basis elements will not change and model runs are obtained at the same grids or scattered locations. In this study we illustrate our technique to a specific horizontal output from the WACCM simulator. The SHT of model outputs can also be extended to time varying processes. As noted by Jones (1963), if a random field on a sphere varies with time, the representation becomes

$$\eta(\mathbf{s}, t) = \sum_{k=0}^{\infty} \sum_{h=-k}^k c_{k,h}(t) \psi_{k,h}(\mathbf{s}),$$

where $c_{k,h}(t)$ being an ordinary one-dimensional stochastic process. The set of all $c_{k,h}(t)$ form an infinite dimensional stochastic process. Theoretically we can formulate model outputs in space–time settings with such representations. Nevertheless, in climate or chemistry–transport simulations, we often encounter not only outputs in time and horizontal resolution, but also in vertical resolution. Therefore extensions to 4 dimensional correlations are needed, but they must maintain the computational tractability.

In our approach the covariance matrix is formulated as a block diagonal structure, we could relax this assumption and then adopt the block composite likelihood approach to accelerate the algorithm (Chang et al., 2015b). Unfortunately, this approach only cover the stationary case (though could be extended). Our approach naturally and efficiently models nonstationarity in space. Furthermore, there are cases where our approach is computationally more efficient than Chang et al. (2015b). Indeed, if m is large, the computational cost is about $O(\sum_{i=1}^B m_i^3)$, where $\sum_{i=1}^B m_i = m$ (depends on number and size of blocks m_i) versus we have $O(N_y^3 r^3)$ that can be lower in many applications.

Finally, it is shown that the basic formulation from Kennedy and O’Hagan (2001) can lead to asymptotically inconsistent calibration (Tuo and Wu, 2015a). To improve the Bayesian calibration, we could follow in the future the L_2 calibration in our setting (Tuo and Wu, 2015b). The benefit would be to yield a better overall match compared to the likelihood distance. This might improve the quality of the calibration further.

References

- Alexander, M. J. and Sato, K. (2015). Gravity wave dynamics and climate: an update from the SPARC gravity wave activity. *SPARC newsletter*, (44):9–13.
- Bayarri, M., Berger, J., Cafeo, J., Garcia-Donato, G., Liu, F., Palomo, J., Parthasarathy, R., Paulo, R., Sacks, J., and Walsh, D. (2007). Computer model validation with functional output. *The Annals of Statistics*, 35(5):1874–1906.
- Beck, J. and Guillas, S. (2016). Sequential design with mutual information for computer experiments (MICE): emulation of a tsunami model. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):739–766.
- Bhat, K. S., Haran, M., and Goes, M. (2010). Computer model calibration with multivariate spatial output: A case study. *Frontiers of Statistical Decision Making and Bayesian Analysis*, pages 168–184.
- Bolin, D. and Lindgren, F. (2011). Spatial models generated by nested stochastic partial differential equations, with an application to global ozone mapping. *Annals of Applied Statistics*, 5(1):523–550.
- Chang, K.-L., Guillas, S., and Fioletov, V. E. (2015a). Spatial mapping of ground-based observations of total ozone. *Atmospheric Measurement Techniques*, 8(10):4487–4505.
- Chang, W., Haran, M., Olson, R., and Keller, K. (2014). Fast dimension-reduced climate model calibration and the effect of data aggregation. *Annals of Applied Statistics*, 8(2):649–673.
- Chang, W., Haran, M., Olson, R., and Keller, K. (2015b). A composite likelihood approach to computer model calibration with high-dimensional spatial data. *Statistica Sinica*, 25:243–259.

- Gneiting, T. (2013). Strictly and non-strictly positive definite functions on spheres. *Bernoulli*, 19(4):1327–1349.
- Higdon, D., Gattiker, J., Williams, B., and Rightley, M. (2008). Computer model calibration using high dimensional output. *Journal of the American Statistical Association*, 103:570–5833.
- Jones, R. H. (1963). Stochastic processes on a sphere. *Annals of Mathematical Statistics*, 34:213–218.
- Jun, M. and Stein, M. L. (2008). Nonstationary covariance models for global data. *Annals of Applied Statistics*, 2:1271–1289.
- Kennedy, M. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63(3):425–464.
- Kirby, M. (2001). *Geometric data analysis : an empirical approach to dimensionality reduction and the study of patterns*. Wiley, New York.
- Lindgren, F. and Rue, H. (2015). Bayesian spatial and spatiotemporal modelling with R-INLA. *Journal of Statistical Software*, 63(19).
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.
- Liu, H., Sassi, F., and Garcia, R. (2009). Error growth in a whole atmosphere climate model. *Journal of the Atmospheric Sciences*, 66(1):173–186.
- McNeall, D., Challenor, P. G., Gattiker, J., and Stone, E. (2013). The potential of an observational data set for calibration of a computationally expensive computer model. *Geoscientific Model Development*, 6:1715–1728.
- Muir, J. and Tkalčić, H. (2015). A method of spherical harmonic analysis in the geosciences via hierarchical Bayesian inference. *Geophysical Journal International*, 203(2):1164–1171.
- Rougier, J. (2008). Efficient emulators for multivariate deterministic functions. *Journal of Computational and Graphical Statistics*, 17(4):827–843.
- Stein, M. L. (2007). Spatial variation of total column ozone on a global scale. *Annals of Applied Statistics*, 1:191–210.
- Tuo, R. and Wu, C. (2015a). A theoretical framework for calibration in computer models: parametrization, estimation and convergence properties. *arXiv preprint arXiv:1508.07155*.
- Tuo, R. and Wu, C. J. (2015b). Efficient calibration for imperfect computer models. *The Annals of Statistics*, 43(6):2331–2352.
- Whittle, P. (1963). Stochastic processes in several dimensions. *Bulletins of the International Statistical Institute*, 40:974–994.
- Williamson, D., Goldstein, M., Allison, L., Blaker, A., Challenor, P., Jackson, L., and Yamazaki, K. (2013). History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. *Climate Dynamics*, 41(7-8):1703–1729.
- Zammit-Mangion, A., Rougier, J., Bamber, J., and Schön, N. (2015). Resolving the Antarctic contribution to sea-level rise: a hierarchical modelling framework. *Environmetrics*, 25(4):245–264.